

The Examiner maintained that the statement in these claims that the addresses in the monitored subset are "expected to receive smaller amounts of the communication traffic than other addresses in the group" was unclear for the following reasons:

- (1) "Why is that?"
- (2) "Is there the purpose of identifying this particular subset as to receive smaller amounts of communication traffic?"
- (3) "Why expecting that aspect?"

Reasons (1) and (2) relate to the motivation for this feature of the invention, and have nothing to do with the clarity or distinctiveness of the claim language. Actually, the very fact that the Examiner was able to ask these pointed questions (asking "why," rather than "what") would tend to indicate that the language of the claim is clear to the Examiner. There is no requirement in 35 U.S.C. 112 that a claim recite the benefits of or motivation for the invention.

Reason (3) would appear to relate to the question of how a person of ordinary skill in the art would identify a network address that is expected to receive a smaller amount of communication traffic than another. This would appear to be not so much an issue of clarity as of enablement. The specification (paragraphs 0061-0062) points out a number of ways that low expected traffic levels may be ascertained, based on either baseline measurements or the type of equipment deployed at the given network addresses. Applicant respectfully submits that identifying a network address that is expected to receive smaller amounts of communication traffic (and even understanding the "why" behind such an expectation) would have been well within the common sense of one of even minimal skill in the art.

Therefore, Applicant respectfully submits that all of the claims in this application meet the requirements of 35 U.S.C. 112.

Claims 1, 4-11, 21, 22, 25, 26, 28-35, 38-45, 55, 56, 59, 60, 62-69, 72-79, 89, 90, 93, 94, 96-103, 105 and 107 were rejected under 35 U.S.C. 102(e) over Lyle (U.S. Patent 6,886,102). Applicant respectfully traverses this rejection.

Lyle describes a system and method for protecting a computer network against denial of service attacks. One element of this system is a sniffer module, which is "used to monitor network traffic at the ports of devices throughout the network..., to identify messages related to a known or suspected attack or to identify messages that satisfy certain pre-configured criteria believed to indicate the likelihood or possibility that an attack is taking place" (col. 7, lines 39-45). Functions performed by the sniffer module may include searching for predefined strings, as well as searching for "other information, clues, or signatures previously associated with attacks," such as messages sent from suspicious source addresses or messages attempting to access a target system "via a service known to be vulnerable" (col. 10, lines 30-43). Sniffers may also monitor switch and router ports "to detect if a particular port is receiving an unusually high number of data packets of any type, a high number of data packets of a particular type, and/or a high number of packets with a certain target destination or recipient address" (col. 10, lines 44-49). A statistics database may be used to determine whether the rate of certain types of messages exceeds a normal level (col. 10, lines 52-55).

Claim 1 recites a method for processing communication traffic that is directed to a group of addresses on a network, based on monitoring traffic that is directed to a subset of the group. The subset of the

group of the addresses is identified such that the addresses in the subset are expected to receive smaller amounts of the communication traffic than other addresses in the group. When the characteristics of the traffic directed to at least one of the addresses in the subset deviates from a baseline - indicating that at least a portion of the traffic is of potentially malicious origin - the traffic that is directed to all of the addresses in the group is filtered so as to remove at least some of the malicious traffic.

As Applicant has pointed out previously, Lyle neither teaches nor suggests any particular criterion for selection of ports or addresses to be monitored by his sniffer. In regard to the "identifying" step of claim 1, the Examiner cited Lyle's Fig. 8 and col. 14, line 56 - col. 15, line 30, in which Lyle's analysis framework determines whether the number of events associated with a network or sub-network "exceeds a baseline incident rate by a prescribed amount" (col. 14, lines 59-62). Elsewhere (col. 8, lines 18-20) Lyle refers to a "baseline incident rate." Neither of these passages has anything to do with identifying or choosing to monitor addresses that are expected to receive smaller amounts of communication traffic, as recited in claim 1. The Examiner has failed to point out even a hint of motivation in Lyle for monitoring low-traffic addresses.

Therefore, claim 1 is believed to be patentable over the cited art. In view of the patentability of claim 1, dependent claims 4-11, 21, 22 and 103 are also believed to be patentable.

Claims 35, 38-45, 55, 56, 69, 72-79, 89, 90, 105 and 107 recite apparatus and computer software products that operate on principles similar to the methods of claims 1, 4-11, 21, 22 and 103. Therefore, claims 35, 38-45, 55, 56, 69, 72-79, 89, 90, 105 and 107 are believed to be

patentable for the reasons explained above with respect to claim 1.

Independent claim 25 recites a method for processing communication traffic in which traffic originating from a group of addresses is monitored in order to detect a pattern that is indicative of a malicious program running on a computer at one (or more) of the addresses. The pattern is detected by determining that the computer has transmitted packets to a large number of different destination addresses. A route of the traffic is traced back to the address so as to identify a location of the computer on which the malicious program is running.

Lyle neither teaches nor suggests applying the sort of detection criterion that is recited in claim 25, based on the large number of different destination addresses to which a computer has transmitted packets. In rejecting claim 25, the Examiner maintained that Lyle teaches this feature in col. 10, lines 19-60, and col. 13, lines 9-21 and 38-55. The passage in col. 10 refers to detection of certain strings, clues, or signatures, as well as detection of a high number of data packets of some type and/or "with a certain target destination or recipient address" (lines 48-49, emphasis added). By contrast, claim 25 recites the opposite criterion: packets directed to many different destination addresses.

The cited passages in col. 13 relate to the manner in which events are handled in order to prevent multiple messages to one destination from masking another message to a different destination (lines 16-18) and to determine whether multiple events should be aggregated into an "existing incident" (lines 37-38). These are internal function of Lyle's system, which have nothing to do with determining that a suspect computer has transmitted packets to a large number of different destination addresses, let alone using the large number of different destination addresses as a pattern indicative of a

malicious program running on the computer, as recited in claim 25.

The Examiner continued this line of misinterpretation of the claim limitations in the rationale for the rejection (page 7, lines 15-18, in the Official Action): "Lyle also discloses that the method of detecting the router ports if a particular ports is receiving an unusually high number of data packets of any type with a certain target destination or recipient address." Claim 25 refers to packets sent to many different destination addresses, and certainly not many packets sent to the same destination address, as the Examiner appears to have interpreted the claim.

Thus, claim 25 is believed to be patentable over the cited art. In view of the patentability of claim 25, dependent claims 26 and 28 are also believed to be patentable.

Claims 59, 60, 62, 93, 94 and 96 recite apparatus and computer software products that operate on principles similar to the methods of claims 25, 26 and 28. Independent claims 59 and 93 have been amended in like manner to claim 25. Therefore, claims 59, 60, 62, 93, 94 and 96 are believed to be patentable for the reasons explained above with respect to claim 25.

Independent claim 29 recites a method in which communication traffic is monitored so as to detect packets indicative of a network communication failure that is characteristic of a worm infection. Upon detecting an increase in the rate of arrival of these packets, the communication traffic is filtered so as to remove at least a portion of the communication traffic that is generated by the worm infection.

The passage in Lyle that the Examiner cited against claim 29 (col. 10, line 53 - col. 11, line 1) relates only to detecting the "level or rate" of "certain types of messages" (lines 55-56), without specifying the types

of messages that are involved. Lyle makes no mention or suggestion of communication failures or how they should be handled, and does not even hint that packets indicative of such failures could be used in filtering worm-generated traffic as required by claim 29.

In the "Response to Arguments" (page 43, first full paragraph, in the present Official Action), the Examiner stated that "Lyle does teach of communication failures or how they should be handled" in col. 14, lines 26-34. This passage, however, relates not to communication failures, but rather to the ways in which Lyle's analysis framework may respond to an event that is indicative of an attack (col. 14, lines 8-12). The response of the framework, as cited by the Examiner, may include sending a message "to stop a malicious flow of network traffic." This response might be considered akin to "filtering the communication traffic," as recited in the last step of claim 29. There is not even the slightest suggestion in Lyle, however, that any sort of response might be invoked upon detecting an increase in the rate of arrival of packets that are indicative of a communication failure in the network. Among the various ways that Lyle enumerates for detecting attacks, none of them have to do with communication failures. In this regard, Applicant agrees with the Examiner's statement in the Response to Arguments that "Lyle... does not even hint that packets indicative of such failures could be used in filtering worm-generated traffic" (page 43, lines 11-12).

Dependent claim 30 depends from claim 29 and adds that "ICMP unreachable" packets are detected in monitoring the communication traffic. The Examiner maintained that Lyle teaches this limitation in col. 9, lines 7-37, but Lyle contains no teaching or suggestion at all of ICMP either in this passage or elsewhere.

Thus, independent claim 29 is believed to be patentable over the cited art. In view of the

patentability of claim 29, dependent claims 30 and 31 are also believed to be patentable.

Claims 63-65 and 97-99 recite apparatus and computer software products that operate on principles similar to the methods of claims 29-31. Therefore, claims 63-65 and 97-99 are believed to be patentable for the reasons explained above with respect to claim 29.

Independent claim 32 recites a method in which communication traffic on a network is monitored so as to detect ill-formed packets. The ill-formed packets are used in determining that at least a portion of the traffic has been generated by a worm infection. The communication traffic is then filtered in order to remove at least the portion of the worm-generated traffic. The meaning of "well formed" in this context and examples of ill-formed packets are described in the specification (paragraphs 0066-0068) and are recited in dependent claims 33 and 34.

Lyle fails to relate in any way to whether packets are well formed or ill formed, and certainly does not suggest that detection of ill-formed packets might be used in determining that a worm infection has occurred. The passage cited by the Examiner against claim 32 (col. 7, lines 9-19) says only that "the sniffers search for data indicating an actual or suspected attack... as described more fully below." Lyle describes a number of ways in which the sniffers may search for such attack-related data (see, for example, col. 10, lines 30-59). None of these ways has anything to do with ill-formation of packets.

In the "Response to Arguments" (starting three lines from the bottom of page 43 in the Official Action), the Examiner maintained that "Lyle discloses that the method of search or determined the suspicious data or suspected an attack, which the same as determined the ill formed packets or worm infection has occurred," citing the

above-mentioned passage in col. 7, along with col. 8, lines 26-39, and col. 10, lines 15-35. The meaning of the Examiner's comment is not clear, although it is possible that the Examiner is implying - incorrectly - that worm-infected packets are by definition ill formed. The passage in col. 8 refers to the use of a policy database to indicate how certain types of events and incidents should be processed (lines 22-24). The passage in col. 10 mentions protocols that may be used for traffic monitoring and states that the sniffer may use string matching to detect attacks (lines 30-34). The appearance of such a string may indeed signal that a worm attack is in progress, but a packet containing such a string may be perfectly well formed.

Ill-formed packets may be indicative of a worm attack, as discovered by the inventors in the present patent application and recited in claim 32, but this does not mean that worm-infected packets are necessarily ill formed. Lyle mentions many possible signs that could be used to detect worm attacks, but he does not even hint that ill-formation of packets could be one of them. He certainly does not refer to the specific types of ill-formation that are recited in claims 33 and 34. Once again, Applicant agrees with the Examiner's conclusion that "Lyle... certainly does not suggest that detection of ill-formed packets might be used in determining that a worm infection has occurred" (page 44, lines 2-4, in the Official Action).

Thus, Applicant respectfully submits that claim 32 is patentable over the cited art, as are claims 33 and 34, which depend from claim 32.

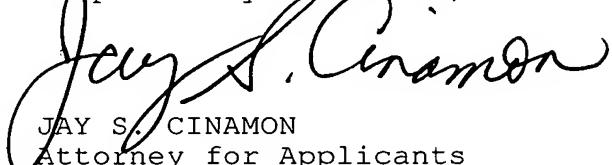
Claims 66-68 and 100-102 recite apparatus and computer software products that operate on principles similar to the methods of claims 32-34. Therefore, claims 66-68 and 100-102 are believed to be patentable for the reasons explained above with respect to claim 32.

Dependent claims 12-20, 23, 24, 46-54, 57, 58, 80-88, 91, 92, 104, 106 and 108 were rejected under 35 U.S.C. 103(a) over Lyle in view of Porras (U.S. Patent 6,321,338) or Trcka (U.S. Patent Application Publication 2001/0039579) or Bartleson et al. (U.S. Patent 6,934,857). In view of the patentability of independent claims 1, 35 and 69, from which these claims depend, dependent claims 12-20, 23, 24, 46-54, 57, 58, 80-88, 91 and 92 are also believed to be patentable.

Notwithstanding the patentability of the independent claims in this application, as explained above, the dependent claims are also believed to recited independently-patentable subject matter. In the interest of brevity, however, Applicant will refrain from arguing the patentability of the dependent claims at present.

Applicant believes the remarks stated above to be fully responsive to all of the grounds of rejection raised by the Examiner. In view of these remarks, all the claims in the present patent application are believed to be in condition for allowance. Prompt notice to this effect is requested.

Respectfully submitted,



JAY S. CINAMON
Attorney for Applicants
Reg. No. 24,156

ABELMAN, FRAYNE & SCHWAB
666 THIRD AVENUE, 10th Fl.
NEW YORK, NEW YORK 10017
(212) 949-9022
(2120) 949-9190